

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE  
BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES

In re Application of:

ODED COHN ET AL.

Serial No. 08/364,334

Filed: 27 DECEMBER 1994

For: CACHE CONTROL SYSTEM

§ Docket No. UK9-93-044

§

§ Examiner: C. CHOW

§

§

§

§ Art Unit: 2318

§

§

#20  
2/28/97  
6-8-97  
RECEIVED

APR 23 1997

BOARD OF PATENT APPEALS  
AND INTERFERENCES

APPEAL BRIEF

Assistant Commissioner of Patents  
Washington, D.C. 20231

Sir:

RECEIVED  
APR 29 1997  
GROUP 2300

This Brief is submitted in triplicate in support of the Appeal in the  
above-identified application.

**CERTIFICATE OF MAILING**  
**37 CFR 1.8(a)**

I hereby certify that this correspondence is being deposited  
with the United States Postal Service as First Class Mail in  
an envelope addressed to: Commissioner of Patents and  
Trademarks, Washington, D. C. 20231 on the date below.

4/18/97  
Date

Andrea L. Bass  
Signature

### **REAL PARTY IN INTEREST**

As reflected in the Assignment recorded February 27, 1995, at Reel 7364, Frame 0664, the present application is assigned to International Business Machines Corporation, the real party of interest.

### **RELATED APPEALS AND INTERFERENCES**

Applicant knows of no appeal or interference related to the present appeal.

### **STATUS OF CLAIMS**

Claims 1-3 and 10-18 stand finally rejected as noted by the Examiner in the final rejected dated November 5, 1996.

### **STATUS OF AMENDMENTS**

No amendment to the claims has been submitted subsequent to the final rejection.

### **SUMMARY OF THE INVENTION**

As set forth in the present specification at page 5, line 15 *et seq.*, the present invention provides a method for operating a data storage apparatus having a movably mounted storage element and a solid-state cache memory, in which the data storage apparatus storage element is not at operating speed when data access has not occurred during a predetermined time period. The present invention accesses data on the storage element in response to a read or write request which cannot be satisfied via access to the cache memory and transfers data between the cache memory and the storage element to maintain consistency of data there between only while the storage element is at operating speed after a read or write request has given rise to an access to the storage element.

As illustrated in **Figure 1** and as described in the present specification at page 6, line 13 *et seq.*, a data processing system is depicted which includes computer **100** which is attached to data storage apparatus **110** via a SCSI (Small

Computer Systems Interface) **101**. Data storage apparatus **110** comprises: magnetic disk storage device **104** which has, in this example, a capacity of 120 Megabytes; cache system **102** which includes a non-volatile solid-state cache memory **106** having a capacity, in this example, of 2 Megabytes and cache control or management logic **108**. Cache control system **108** includes a suitable microprocessor together with an associated control code to enable the microprocessor to perform the various functions described herein and cache management tables for storing control data for the cache memory and data storage devices.

Disk storage device **104** has three possible states: standby, idle and active. In the active state, the disk is rotating at operating speed and Read/Write and Seek operation may be performed. The typical power required in this state can be between 1.5W and 2.5W. In the idle state, the disk is still at operating speed and the power consumption is 0.9W. There is no power overhead when going from the idle state to the active state. The standby state is the state in which the disk is stationary and has the lowest possible power consumption, typically 0.2W. Switching from the standby state to the active state, necessitates the disk to be brought to operating speed. This spin-up typically requires a power of up to 4W over the acceleration time. The acceleration time is relatively long compared to an I/O operation.

Data within cache memory **106** corresponding to data stored within the disk can be in one of three states: "Invalid", a sector which is designated as invalid does not contain meaningful data; "Consistent", a sector which is designated consistent contains the same information as that stored within the corresponding disk sector; and, "New", a sector which is designated new contains information fresher than that stored on the corresponding sector of the disk.

As described at line 24, page 12, the cache control system has a cache replacement mechanism which performs two main tasks: destaging, i.e., transferring data from cache to disk and replacement. The cache control system

executes destage operations only when the disk is spinning anyway. Only new sectors are destaged. Each sector which has been destaged is marked as consistent.

In accordance with an important feature of the present invention, and as described at page 13, line 15 in the present specification, each time a read or write request has given rise to an access to the disk storage device, a destage takes place for a certain period of time if there are new sectors in the cache. The maximum period for which destaging takes place on each occasion is determined by allocating a certain energy budget to that destaging. For example, if keeping the disk in an active state for  $t$  time units takes the same energy as required to spin up the disk, then, destaging will occur during at  $+ b$  units of time, where  $a$  and  $b$  are system-dependent tunable parameters.

As depicted in **Figure 4** and as described in the present specification at page 14, line 16, any read request is first examined to see if the request is for more than a predetermined number of sectors. This threshold will be dependent upon the size of the cache and the characteristics of the disk. Large requests are satisfied by transferring the data directly from the disk to the main memory of the computer. Since the disk is in a read/write state, and the disk must be accelerated to its operating speed, any required destaging from the cache to the disk is also performed for a period of time equal to the destaging interval.

### **ISSUES**

1. Is the Examiner's rejection of Claims 1, 2, 12 and 18 under 35 U.S.C. § 103 as being unpatentable over *Yamazaki*, Japanese Patent Application Laid-Open 4-205852 well founded?

2. Is the Examiner's rejection of Claim 3 under 35 U.S.C. § 103 as being unpatentable over *Yamazaki* and further in view of *Hanson et al*, United States Patent No. 4,433,374 well founded?

3. Is the Examiner's rejection of Claims 10 and 14 under 35 U.S.C. § 103 as being unpatentable over *Yamazaki* well founded?

4. Is the Examiner's rejection of Claims 11, 13 and 15 under 35 U.S.C. § 103 as being unpatentable over *Yamazaki* in view of "art, knowledge" well founded?

5. Is the Examiner's rejection of Claim 16 under 35 U.S.C. § 103 as being unpatentable over *Yamazaki* and further in view of *Noya et al*, United States Patent No. 5,420,983 well founded?

6. Is the Examiner's rejection of Claim 17 under 35 U.S.C. § 103 as being unpatentable over *Yamazaki* and *Noya et al* in view of "art, knowledge" well founded?

#### **GROUPING OF THE CLAIMS**

For purposes of this appeal, Claims 1-3 and 10-18 stand or fall together as a single group.

#### **ARGUMENT**

The Examiner has rejected Claims 1, 2, 12 and 18 under 35 U.S.C. § 103 as being unpatentable over *Yamazaki*, Japanese Patent Application Laid-Open 4-205852. That rejection is not well founded and it should be reversed.

The Examiner takes the position that *Yamazaki* teaches a data storage system having a solid state cache memory, a storage element having a moving part, a cache control system and a means to access data stored on the storage element if a read/write request cannot be satisfied via access to the cache memory and means for accessing data stored within the cache if the read/write request from the computer can be satisfied via an access to the cache.

The Examiner takes the position that *Yamazaki* also "inherently teaches a means for designating selected data within the cache as new data in response to a write from the computer which updates data within the cache because *Yamazaki* manages write requests, which would necessitate a means of tracking dirty blocks in order to maintain cache coherency." Further, the Examiner takes the position that *Yamazaki* differs from the claimed invention by "not specifically teaching a cache replacement mechanism which flushes dirty entries when another disk transition requires access to the disk itself."

Ignoring for the moment the fact that Applicant's claimed invention has nothing to do with "flushing dirty entries", Applicant would urge the Board to review **Figure 1** of *Yamazaki*. **Figure 1** discloses that a request for data from the host which results in a cache hit causes the data to be exchanged between the cache memory and the host. As a result of a cache miss, the disk spindle motor is turned on and data is exchanged between the disk and the host via the cache memory and the disk spindle motor is thereafter turned off.

The flow chart within **Figure 3** of *Yamazaki* is similar but includes the additional feature of determining whether or not the hit rate for data required by the host within the cache memory is high or low. Thus, as a result of a miss, the hit rate is determined and if the hit rate is high, the disk spindle motor is turned on and the data is exchanged between the disk and host via the cache memory and thereafter the disk spindle motor is turned off. However, if the hit rate is low, the spindle motor is switched on, if it is previously off and data is exchanged between the disk and host via the cache memory.

Applicant has carefully examined *Yamazaki* in full detail and fails to find any recitation within *Yamazaki* of writing data to either the disk memory or the cache. As the reference fails to show or suggest in any way the writing of memory to the cache system or disk, Applicant urges that this reference cannot be said to show or suggest the system of Claim 1, for example, which positively recites the designation of selected data within the cache memory as "new data in response

to a write request from the computer which updates within the cache memory ..." (*Emphasis added*) and a cache replacement mechanism which thereafter transfers such "new" data from the cache memory to the storage element "when the storage element is at operating speed as a result of a read or write request which requires access to the storage element," where "new" data is a term defined in Applicant's specification. *Yamazaki* is absolutely silent on the subject of modified data within the cache and how that data is transferred to the disk storage element and absent some suggestion within the reference, Applicant urges that the Examiner has failed to make a *prima facie* case of obviousness in this rejection. Further, Applicant urges that one having ordinary skill in the art would not find a suggestion for storing new data within the cache and thereafter destaging data from the cache to the disk storage element only when the disk storage element has been brought to operating speed as a result of another read or write operation in the manner set forth within the present claim. Consequently, Applicant urges that the Examiner's rejection of Claims 1, 2, 12 and 18 over the *Yamazaki* reference is not well founded.

Returning now to the Examiner's characterization of *Yamazaki* as deferring from the claimed invention by not specifically teaching a cache replacement mechanism "which flushes dirty entries when another disk transaction requires access to the disk itself ...", Applicant would urge the Board to consider that there is a substantial difference between the "flushing" of a "dirty" entry from the cache and the transfer of updated or "new" data from cache to the disk storage device in the manner which is expressly defined within the present specification and set forth within the present claims. The Examiner's gross mischaracterization of the present invention is believed by the Applicant to form the basis for the Examiner's misapprehension that *Yamazaki* is suggestive of the present invention and reversal of this rejection is respectfully requested.

The Examiner has rejected Claim 3 under 35 U.S.C. § 103 as being unpatentable over *Yamazaki* and further in view of *Hanson et al*, United States Patent No. 4,433,374. That rejection is also not well founded and it should be reversed.

The Examiner cites *Hanson et al* for its teaching that a cache/disk subsystem may be provided in which long data transfers may be written directly to the disk subsystem, bypassing the cache unit. A careful review of *Hanson et al* reveals not the slightest suggestion therein of annotating data within the cache as new data and thereafter transferring that data from the cache to the disk storage element when the disk storage element "is at operating speed as a result of a read or write request which requires an access to the storage element." Consequently, Applicant urges that no combination of *Hanson et al* and *Yamazaki* can be said to show or suggest the invention set forth within Claim 3 and reversal of this rejection is respectfully requested.

The Examiner has rejected Claims 10 and 14 under 35 U.S.C. § 103 as being unpatentable over *Yamazaki*. That rejection is also not believed to be well founded and reversal of this rejection is also respectfully requested.

Claims 10 and 14 depend from Claim 1 and thus, incorporate the features set forth therein. Specifically, the annotation of data within the cache memory as new data in response to a write request from the computer which updates the data within the cache memory and the subsequent transfer of that data from the cache memory to the storage element when the storage element is at operating speed as a result of a read or write request which requires access to the storage element. Consequently, Applicant urges that the Examiner's rejection of Claims 10 and 14 under 35 U.S.C. § 103 as being unpatentable over *Yamazaki* is not well founded and reversal of that rejection is respectfully requested.

The Examiner has also rejected Claims 11, 13 and 15 under 35 U.S.C. § 103 as unpatentable over *Yamazaki* and further, in view of "art common



knowledge". Claims 11, 13 and 15 depend, either directly or indirectly from Claim 1 and thus incorporate the features of Claim 1. As *Yamazaki* fails to show or suggest in any way the novel power saving method and system of the present invention, whereby data is updated to the cache memory and designated as new data thereafter transferred from the cache memory to the storage element when the storage element is at operating speed as a result of a read or write request which requires access to the storage element, the Applicant urges the Board to reverse the Examiner's rejection of Claims 11, 13 and 15. No combination of *Yamazaki* and common prior art knowledge can be said to show or suggest this novel power saving method and system.

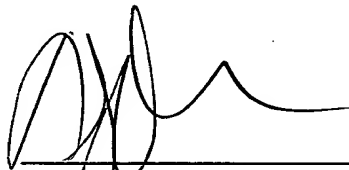
The Examiner has also rejected Claim 16 under 35 U.S.C. § 103 as being unpatentable over *Yamazaki* as applied to Claim 1 and further in view of *Noya et al*, United States Patent No. 5,420,983. That rejection is not well founded and it should be reversed. The Examiner cites *Noya et al* for its teaching of a non-volatile memory for use as a disk write cache in combination with a disk subsystem; however, no combination of *Noya et al* and *Yamazaki* shows or suggests in any way the system set forth within Claim 1, from which Claim 16 depends. For the reasons set forth above, with respect to Claim 1, this rejection is also urged as not well founded and reversal of the Examiner's rejection is respectfully requested.

Finally, the Examiner has rejected Claim 17 under 35 U.S.C. § 103 as being unpatentable over *Yamazaki* and *Noya et al* as applied to Claim 16 and further in view of "art common knowledge". The Examiner cites this combination of references for a showing that the disk subsystem be utilized with a portable or personal computer; however, Claim 17 depends, indirectly, from Claim 1 and thus incorporates all features of Claim 1 as set forth and argued above. No combination of *Yamazaki* and *Noya et al* and the common knowledge cited by the Examiner can be said to show or suggest the novel techniques set forth herein wherein the writing of data to the cache and designating of data as "new" data occurs and wherein that "new" data is thereafter destaged or transferred from the cache

memory to the data storage subsystem as a result of a read or write request which has required the data storage subsystem be brought to operating speed. In view of the above, Applicant urges the Board to reverse each of the Examiner's rejections with respect to Claims 1-3 and 10-18.

Please charge the fee of **\$300.00** for submission of a Brief in Support of Appeal to Deposit Account No. **09-0465**. No additional fee is seen to be required; however, in the event an additional fee is required, please charge that fee to Deposit Account No. **09-0465**. No extension of time is believed to be required; however, in the event an extension of time is required, please consider that extension requested and please charge any associated fee therefore to the above-identified deposit account.

Respectfully submitted,

A handwritten signature in black ink, appearing to read 'Andrew J. Dillon', is written over a horizontal line.

Andrew J. Dillon  
*Registration No. 29,634*  
FELSMAN, BRADLEY, GUNTER  
& DILLON, LLP  
Suite 350 Arboretum Point  
9505 Arboretum Boulevard  
Austin, Texas 78759  
512.343.6116

## APPENDIX

1 1. A data storage system for fulfilling read and write requests from a computer  
2 comprising:

3 a solid-state cache memory;

4 a storage element with at least one moving part wherein the storage element  
5 has an operating, and a non fully operating mode when data access has not  
6 occurred for a predetermined time period;

7 a cache control system responsive to read and write requests from the  
8 computer, said cache control system further comprising,

9 means to access data stored on the storage element if a read or write  
10 request cannot be satisfied via access to the cache memory;

11 means for accessing data stored within the cache memory if a read or write  
12 request from the computer can be satisfied via access to the cache memory;

13 means for designating selected data within said cache memory as new data  
14 in response to a write request from the computer which updates data within the  
15 cache memory; and

16 a cache replacement mechanism for transferring new data from the cache  
17 memory to the storage element to maintain consistency of data between data  
18 stored in the cache memory and data stored in the storage element, wherein the  
19 cache replacement mechanism performs data transfers between the cache memory  
20 and the storage element when the storage element is at operating speed as a result  
21 of a read or write request which requires an access to the storage element.

1 2. The data storage apparatus of Claim 1 wherein the cache replacement  
2 mechanism carries out transfers of new data from the cache memory to the storage  
3 element for a predetermined time each time a read or write request has given rise  
4 to an access to the storage element.

1 3. The data storage apparatus of Claim 1 wherein the cache control system further  
2 comprises means for accessing the data on the storage element without modifying  
3 the data in the cache memory if the read or write request involves more than a  
4 predetermined amount of data.

1 10. The data storage apparatus of Claim 1 wherein the storage element with at  
2 least one moving part is an optical drive.

1 11. The data storage apparatus of Claim 10 in the form of a disk storage  
2 subsystem for use with a personal or portable computer.

1 12. The data storage apparatus of Claim 1 wherein the storage element with at  
2 least one moving part is an magnetic disk drive.

1 13. The data storage apparatus of Claim 12 in the form of a disk storage  
2 subsystem for use with a personal or portable computer.

1 14. The data storage apparatus of Claim 1 wherein the storage element with at  
2 least one moving part is an magnetoptical disk drive.

1 15. The data storage apparatus of Claim 14 in the form of a disk storage  
2 subsystem for use with a personal or portable computer.

1 16. The data storage apparatus of Claim 1 wherein the solid state cache memory  
2 is non-volatile.

1 17. The data storage apparatus of Claim 16 in the form of a disk storage  
2 subsystem for use with a personal or portable computer.

1 18. A method for operating data storage apparatus having a storage element  
2 including at least one moveable part and a solid-state cache memory in which the  
3 storage element is in a non-fully operational mode when data access has not  
4 occurred during a predetermined time period, the method comprising the steps of:

5 accessing data on the storage element in response to a read or write request  
6 which cannot be satisfied via access to the cache memory;

7 accessing data stored within the cache memory if a read or write request  
8 from the computer can be satisfied via access to the cache memory;

9 designated and selected data within said cache memory as new data in  
10 response to a write request from the computer which updates data within the  
11 cache memory; and

12 transferring new data from the cache memory to the storage element to  
13 maintain consistency of data therebetween, said data transfers between the cache  
14 memory and the storage element being performed when the storage element is fully  
15 operational as a result of a read or write request has given rise to an access to the  
16 storage element.